

(19)



Europäisches Patentamt
European Patent Office
Office européen des brevets

(11) Publication number:

0 336 586
A2

(12)

EUROPEAN PATENT APPLICATION

(21) Application number: 89302584.1

(51) Int. Cl.4: G06F 15/40

(22) Date of filing: 16.03.89

(30) Priority: 08.04.88 US 179316

(43) Date of publication of application:
11.10.89 Bulletin 89/41(84) Designated Contracting States:
BE CH DE ES FR GB IT LI NL SE(71) Applicant: International Business Machines
Corporation
Old Orchard Road
Armonk, N.Y. 10504(US)(72) Inventor: Alexander, Virginia Lee
10506 Mourning Dove Drive
Austin Texas 78750(US)
Inventor: Gaudet, James Louis
1500 Creek Hollow
Austin Texas 78754(US)
Inventor: Jordan II, Lloyd Eugene
13505 Bayswater Garden
Austin Texas 78729(US)
Inventor: Hernandez, Raymond
13304 Council Bluff Drive
Austin Texas 78727(US)(74) Representative: Grant, Iain Murray
IBM United Kingdom Limited Intellectual
Property Department Hursley Park
Winchester Hampshire SO21 2JN(GB)

(54) Data storage, retrieval and transmission in computer systems.

(57) In order to provide for the efficient storing, accessing and/or transmission of a relational data table of a computer system, the table is stored as a plurality of records, arranged successively in series, said records comprising a plurality of data ("D") records of variable length and including items of data for a particular row arranged serially in succession, thereby forming a plurality of data rows and columns with each item of data in each D record arranged in a separate column; and a plurality of column descriptor ("C") records, each C record being associated with one table column and specifying:

(i) the identification of prescribed attributes of the column associated with the respective C record, such as column name, data type, codepage, etc.;

(ii) the identification of the associated D record in a table row which contains the items of data to be located in the associated column;

(iii) the maximum number of characters forming an item of data in the associated column if the number is not implied by the data type therein; and

(iv) the position of the associated column in associated D record,
whereby null or varying length data items which terminate a D record are compressible.

EP 0 336 586 A2

COLUMN NAME	NAME	STATE	REGION	POPULATION
DATA TYPE # CHARACTERS NULLABLE ?	VARCHAR 30 REQUIRED	CHAR 2 REQUIRED	VARCHAR 30 NULLABLE	INTEGER 8 REQUIRED
D RECORD NUMBER	1	2	2	3
D RECORD POSITION	BEGINNING (OF DATA PORTION OF D RECORD)	BEGINNING	BEGINNING PLUS LENGTH OF STATE COLUMN	BEGINNING

FIG. 4

DATA STORAGE, RETRIEVAL AND TRANSMISSION IN COMPUTER SYSTEMS

The present invention relates to data storage, retrieval and transmission in computer systems.

Database manager programs are known which facilitate the storage and retrieval of data stored on a record medium, such as a Winchester hard disk or a 5-1/4 inch diskette, in a digital computer system. Particularly as the size of computers and their associated peripheral storage devices is reduced, it is desirable to make efficient use of the available storage area.

The data format in a database normally includes a plurality of data fields of various sizes. The size of each field is set equal to the number of bytes in an item of data for a fixed length item, or equal to the maximum length of an item of data for a variable length item. In cases of variable length data items, or where an item is not always present for a particular field, the data can be compressed by reducing the size of the fields to accommodate only the items of data actually present.

Thus, for example, in an employee record which comprises successive fields for the employee's name, social security number, birthdate, etc., the name field will be of variable length whereas the social security number and birthdate fields will be of fixed length. Of the latter two fields, it may always be necessary to enter a social security number in its respective field; however, it may not be necessary to include the birthdate. Consequently, the data in these respective fields may take the following form:

JOHN E DOE/987654321/012345

MARY J POPPINS/123456789/000000

The first or name field must be long enough to accommodate a name of any length; the second field must be exactly nine characters long and the third field six characters long. In some cases, no data is entered into the birthdate field, as indicated in the example above by the O's.

When these items of data arranged in these fields are stored, it is desirable to store only the actual useful data. To accomplish this, the database manager must vary the size of the fields in storage so that they are no larger than the actual data requires. Depending upon the data format in storage, this may be a simple, or a very difficult task.

One data format which does not lend itself to straightforward data compression is a so-called "relational database table". A "relational table" is a two dimensional array of data items; that is, an array of data items which form rows and columns. Each row in the table comprises one or more records of information. Each data record comprises one or more fields in which are located the items of data. The fields may have different attributes, such as a social security number, birthdate, etc., or a group of fields may have the same attribute, such as a group of statistical samples for a given parameter. A column in the table is the same field in all of the records. All data items for a column must be of the identical data type (integer, floating point number, fixed length character, variable length graphic, etc.).

Consequently, any specific data item in the table belongs to exactly one row and one column, and can be located as such.

A typical use of a relational table is to describe a collection of homogeneous data objects. A row represents a single instance of an object and the columns are descriptive data attributes about the object. For example, a table of data describing employees could have a row defined for each employee with column values for name, social security number, birthdate, etc.

A "column" in a relational table may therefore be defined as one unit of the vertical dimension of the table. A column normally has a specific data attribute which applies to all data items having the column location. A column, which is analogous to a data "field", is applicable to all rows in the table.

A "row" in a relational table may be defined as one unit of the horizontal dimension of the table. A row contains data items from each column in the table although data items may be missing from some of the row and column locations. All columns are considered to be in the same order for all rows. A row is thus analogous to a "record" in a non-relational database.

There is no presumption of the order of rows within a table. An index can be created by the user to logically order the rows of a table in different ways, each index thus defining a separate ordering. The term, "index" may thus be defined as a collection of the columns in a table with ascending or descending order directives. Once defined, an index is internally maintained by the database manager. It is never referenced in user requests for manipulating or retrieving data, but is used internally whenever an access is determined to be faster using an existing index than through a sequential file search.

An index can be created at any time and dropped at any time. There is no requirement that any index exists for a table.

A "base table" is a relational table that is physically stored on a record medium as real data. A "view" is a relational table that does not exist in physical storage but is derived from one or more base tables.

Given this type of data, a "database" may be defined as a collection of relational tables, catalogue tables and recovery logs. A database may, for example, be restricted to being within a single OS/2™ file system. This restriction ensures all catalogue and recovery functions can be performed on any data under its control at any time. Databases may be stored on hard files or on diskettes. Multiple databases can be placed on a single OS/2™ file system.

Another term which is necessary to understand for the present invention is the so-called "codepage" environment of a database. A "codepage" may be defined as a translation table or mapping of the set of bytes (of which there are 256) into character representations produced by a printer, display, or other input/output device. Well known codepages include the ASCII (American National Standard Code for Information Interchange) IBM National and International codepages 437 and 850, respectively, as well as the EBCDIC character set (Extended Binary-coded Decimal Interchange Code).

In interchanging data between one database and another, it is necessary to know, and keep track of the codepage of the data so that data integrity may be maintained from one codepage to another, if required. However, the interchange method should also allow the structure that exists on the source database to be created on the target database. This means that the table name, column names, column types and indices must be carried as part of the interchange.

It is an object of the present invention to enable a relational database table to be stored and accessed and/or transmitted in an efficient manner in and by a computer system.

The technique provided enables a relational data table to be formatted in a manner which allows file compression of items of data which are null or shorter than a given maximum length, and therefore provides for efficient storage and transmission while maintaining efficient accessibility. Moreover, use of the technique provides a file format which facilitates the exchange of the contents of a relational data table between databases.

From a first aspect of the invention there is provided a relational data table for a computer system comprising a plurality of items of data arranged logically at the intersections of a plurality of table rows and a plurality of table columns, the table being formed by a plurality of records, such records comprising;

(1) a group of one or more data ("D") records for each table row, each D record being of variable length and including one or more items of data for that row arranged serially in succession and terminating with a data character which may be null or of a variable length, the data items for the rows being ordered in a consistent manner within their respective groups of D records,

(2) a plurality of column descriptor ("C") records, each C record being associated with one table column and identifying the location of the record of each group which contains the data item for the column, and the beginning position of the associated data item column in the associated D record; whereby null or varying length data items which terminate a D record are compressible.

In accordance with a second aspect of the invention there is provided a method of storing for accessing and/or transmitting, a relational data table of a digital computer system comprising a plurality of items of data logically arranged at the intersections of a plurality of table rows and a plurality of table columns, the method including the step of storing, on a record medium, the table as a plurality of records arranged successively in series, such records comprising;

(1) a group of one or more data ("D") records for each table row, each D record being of variable length and including one or more items of data for that row arranged serially in succession and terminating with a data character which may be null or of a variable length, the data items for the rows being ordered in a consistent manner within their respective groups of D records,

(2) a plurality of column descriptor ("C") records, each C record being associated with one table column and identifying the location of the record of each group which contains the data item for the column, and the beginning position of the associated data item column in the associated D record; whereby null or varying length data items which terminate a D record are compressible.

Preferably, each C record includes an indication of the maximum number of characters forming an item of data in the associated column, if that number is not implied by the type of data in the data item.

Preferably, also, all of the column descriptor ("C") records precede the data ("D") records in the series arrangement of the records during transmission and storage/retrieval.

The column descriptor ("C") records are preferred to have data ("D") record identifiers and positions in ascending order; whereas the column descriptor records themselves may be arranged in any order.

Two fields may be provided within each of the column descriptor ("C") records for locating the data for the column within a row's worth of data; namely, the data record identifier, which specifies the particular data ("D") record (of possibly several data records comprising a row of data) which contains the data for this column, and the data record position which is used to locate the data for the column within the previously selected data record which is part of a row of data. This record position thus represents the

absolute displacement of the item of data for the respective column within the particular data record.

The file format also makes provision in a C record for the identification of the associated column (column name) as well as the attributes of this column, such as data type, data precision, data nullability and codepage environment. All these attributes of each column are therefore stored only once in a table. In addition, the format is easily expandable for additional column attributes.

No data loss problems are attributable to the file format since it does not force any conversions of database data to occur when transferring data from one database to another. The absence of forced conversions tends to minimise processing time when the data exchange is between like systems.

In general, therefore, using the file format described above, any column is compressible if it is the last column within its data record. This can be accomplished by starting a new data record after any varying length or nullable data field. The beginning of this field is identified by the column descriptor record. The end of this field is preferably identified by a preceding null indicator and/or length indicator in the 'D' record.

Preferably, only those nullable or varying length fields are compressed which have a maximum length longer than an arbitrary constant. For example, such constant may be made equal to the longest length of any supported numeric value. Variable length fields identified as potentially very long (e.g. by data types designated as "LONG VARCHAR" and "LONG VARGRAPHIC") may each occupy their own exclusive data record.

In addition to the column descriptor ("C") and data ("D") records, the relational table may contain other records such as a header ("H") record, a table ("T") record, index ("AI") records and an end ("AE") record. The H, T and AI records may each specify several items of information as will be discussed in detail hereinbelow. The AE record preferably specifies only the end of the table.

A record length field can be added to the beginning of all records (H records, T records, C records, D records, AI records and AE records). The C records can precede the D records in the series arrangement of records on the record medium.

Single byte and double byte codepage fields are added to the H record which begins the table file. Single byte and double byte codepage fields are also included in the C records. The codepage fields in the H record specify the codepage(s) of character data in the table with two possible exceptions:

- (1) Column data within data portions of data records; and
- (2) Data inside additional application-specific records.

The codepage fields in the C records of character (and graphic) columns specify the codepage(s) of the data for the respective column. A character column of "bit data" may have zeros for its single and double byte codepages to signify that the bit data will not be translated into characters when printed.

These codepage identification fields allow programs responsible for migration of database data between a host computer and a personal computer to identify which fields are to have character translations applied. They also allow the user a choice of options when migrating data from one database to another; for example, they permit the user to accept or reject data coming from a different codepage environment.

The present invention will be described further by way of example with reference to an embodiment thereof as illustrated in the accompanying drawings, in which:

Fig. 1 is a block diagram showing a small network comprising a host computer and two personal computers;

Fig. 2 is a block diagram of a personal computer system using an operating system that supports a database manager and two separate databases;

Fig. 3 is an example of data which may be arranged in a relational table;

Fig. 4 is a specific example of a relational table, containing columns of different data types and data records of different length;

Fig. 5 is a chart showing the arrangement of records in the PC IXF file; and

Figs. 6A and 6B (collectively) are a flow chart for writing a PC/IXF file.

Fig. 1 shows the general computer hardware environment of the present invention. The present invention may be implemented on a single personal computer, two or more personal computers connected together to exchange information or, as shown in Fig. 1, two or more personal computer connected together and to a host computer such as an IBM 370 system. Fig. 1 illustrates a host computer 10 which operates with a large disk storage file 12. A first personal computer 14 and a second personal computer 16 are connected to the host computer 10 via multiple lines 18 and 20, respectively, which carry data and control information. The two personal computers are also connected together via data and control lines 22.

The database environment of the present invention can take the form illustrated in Fig. 2. In this case, the PS/2™ personal computer operates with an OS/2 Extended Edition™ Database Manager adaptation of

the Integration Exchange Format (IXF) data interchange software that enables the exchange of relational table structures and data. The character data is stored in one or more databases in a specific codepage environment such as IBM's ASCII codepage 437. Numeric and date/time data are stored internally in a format that the underlying operating system and/or hardware supports.

Specifically, the Database Manager is a database management system (hardware and software) that supports the relational database model in which all data is viewed as a collection of tables. The Database Manager provides a relational command processor called "Database Services" described hereinbelow; a generalised query system for locating data; a system for import and export of data from and to another computer system; and a system for backup and restoration of an individual relational database table, and for table maintenance.

The Database Services is the relational command processor of the Database Manager. It serves a large number of functions which include: a system for storage access; structure query language (SQL) statement processing; database management; lock management; concurrency control; write-ahead logging; recovery services, row-level locking granularity; data recovery in the event of application, system, or record medium failure; and security control.

The PC/IXF architecture allows its Database Services to export a database without need to anticipate the requirements and idiosyncrasies of a receiving (hardware and software) database system. Likewise, a database system importing a PC/IXF file need only understand the PC/IXF architecture; the characteristics of the system which exported the file are of no concern. The PC/IXF file architecture maintains the independence of both the exporting and importing database systems.

The IXF architecture, as a generic relational database exchange format, preferably supports a large set of relational data types, including some data types which perhaps are not supported by a specific relational database product. The PC/IXF file format preserves this flexibility; for example, the PC/IXF architecture supports both single-byte and double-byte string data types. Not all database systems will support all PC/IXF data types; however, even restricted database system implementations will provide for detection and disposition of non-supported data types during import.

Although the present invention is described hereinafter with specific reference to the PC/IXF file, the teachings of the invention are not intended to be limited to a PC environment exclusively. In particular, the invention is not machine-dependent to the PC but can be adapted for use with other computers such as the PS/2™ and RT™ computers of the IBM Corporation, mid-range computers, and the like.

Likewise, while the present invention is described with reference to the OS/2™ operating system, the invention is not so limited and can be adapted for use with any well known operating system.

OS/2™, OS/2 Extended Edition™, OS/2 EETM™, PS/2™ and RT™ are trademarks of IBM Corporation.

Fig. 3 illustrates a typical relational database table. A file which contains such a table uses the terminology table/row/column as opposed to file/record/field or relation/tuple/attribute. As may be seen, the table name is "Cities". The table contains a plurality of horizontal rows, one for each city, each row comprising a plurality of information fields. These information fields are arranged in vertical columns, one for each information field.

In practice, the data records contained in each row of a relational table are selected in a manner shown in Fig. 4. This figure illustrates a possible scheme for allocating D records in the table of Fig. 3. At the top of each column is shown the column name and the type of data contained in the column. The data types are selected from a prescribed set of data types which include, for example:

INTEGER: A four byte integer which represents a whole number between -2,147,483,648 and +2,147,483,647.

VARCHAR: A varying length character string. The maximum length for the string (in bytes) may not exceed 254 bytes. The string is in the codepage prescribed for this column.

CHAR: A fixed length character string with a maximum string length of 254 bytes. The string is in the codepage prescribed for this column.

LONGVARCHAR: A varying length character string which may not exceed 32,767 bytes. This string is in the codepage indicated for this column. The string itself is preceded by a current length indicator which is a two byte integer specifying the length of the string in bytes.

SMALLINT: A two byte integer which represents a whole number between -32,768 and +32,767.

FLOATING POINT: Either a long (eight byte) or short (four byte) floating point number in IEEE format.

GRAPHIC: A fixed length string of double-byte characters which may not exceed 127. The string is in the codepage specified for this column.

VARGRAPHIC: A varying length string of double byte characters, the maximum number of which may not exceed 127. The codepage for this VARGRAPHIC data is as specified for the column.

To permit compression of the stored data, the row data is divided up into separate data ("D") records

(three in this case), numbered 1, 2 and 3. Each D record terminates at either a varying length field (column) or a field (column) which may contain no data (nulls). In this way, the data records in a particular row need be no longer than the necessary size for the actual data contained in their last field.

As explained above, the invention provides a file format including column descriptor ("C") records which specify the associated D record identification number (1, 2, 3, etc.) and the associated D record position. The D record position is stored as a byte displacement within the data area portion of the D record for the beginning of the column.

In general, a PC/IXF file to be written or read by its Database Services, or output by Database Services, consists of an unbroken sequence of variable length records. The file will have the following types of records in the order given:

One header ("H") record of;

One table ("T") record;

Multiple column descriptor ("C") records (one record for each column of data in the table);

Multiple data ("D") records (each row in the table being represented by one or more "D" records); and

Zero, one, or Multiple index ("AI") records (each index being defined from the table); and One end ("AE") record.

A PC/IXF file may also contain application specific ("A") records anywhere after the "H" record. These "A" records are allowed in PC/IXF files to enable an application to include additional data, not defined by the PC/IXF format, in a PC/IXF file. For example, "A" records may be used to store indexes of tables and/or to allow multi-diskette storage of a single IXF file. "A" records are to be ignored by any program reading a PC/IXF file which does not have particular knowledge about the data format and content implied by the application identifier in the "A" record.

In the following table, each specific record is described as a sequence of fields, each of which is given a FIELD NAME. For each record, these fields are of a defined length and appear in the order shown.

HEADER RECORD	
FIELD NAME	FIELD LENGTH (Bytes)
record length	06
record type = "H"	01
IXF identifier	03
IXF version	04
software product	12
date written	08
time written	06
heading record count	05
single byte codepage	05
double byte codepage	05
reserved	02

TABLE RECORD	
FIELD NAME	FIELD LENGTH (Bytes)
record length	06
record type = "T"	01
name length	02
name of data (table name)	18
qualifier	08
data source	12
data convention = "C"	01
data format = "M"	01
machine format = "PC"	05
data location = "I"	01
"C" record count	05
reserved	02
data description	30

COLUMN DESCRIPTOR RECORD	
FIELD NAME	FIELD LENGTH (Bytes)
record length	06
record type = "C"	01
column name length	02
column name	18
column allows nulls	01
column selected flag	01
key column flag	01
data class	01
data type	03
single byte codepage	05
double byte codepage	05
column data length	05
"D" record identifier	03
column position	06
column description	30
number of dimensions	02
size of each dimension	varying

DATA RECORD	
FIELD NAME	FIELD LENGTH (Bytes)
record length	06
record type = "D"	01
"D" record identifier	03
reserved	04
columnar data*	varying

*including null indicator values and lengths of variable fields.

APPLICATION RECORD	
FIELD NAME	FIELD LENGTH (Bytes)
record length	06
record type = "A"	01
application identifier	12
Applic. specific data	varying

These records are arranged in the file in the order illustrated in Fig. 5. The application record can be located between any other two record types.

Figs. 6A and 6B illustrate how a database utility operates to write a PC/IXF file on a record medium. The process for reading data from the record medium is similar to the procedure for writing.

The first step in storing a table is to write the H record, which includes the codepage environment of the source database. Thereafter, the T record is generated and written onto the record medium. Following the T record, the C records are written one at a time. As explained above, these C records are each associated with one table column and specify:

- (1) the identification and prescribed attributes, such as data type, of the column associated with the respective C record;
- (2) the identification of the associated D record in a table row which contains the item of data to be located in the associated column;
- (3) the maximum number of characters forming an item of data in the associated column if the number is not implied by the data type; and
- (4) the position of the associated column in the associated D record.

After all the C records have been written, the program proceeds to write the individual D records for each successive row in the table. The D records are of variable length as the data within the table requires. Preferably, the last field in each D record is either the last column in a row or a column which contains either a variable length data item or a data item which can be nulled.

As may be seen in Fig. 6B, the algorithm first asks if all the rows of data have been written and, if not, whether all the D records in the current row have been considered. If there are no more D records and no more rows to write, the program may write same application specific information such as indexes and/or an ending records, and then the end ("AE") record at the end of the table. If more D records are to be written, the program looks for the last column in the current D record and then computes the D record size from the column data stored in a buffer. This is a variable based on the end of the last significant data in the buffer. Once the size is computed, the D record is written onto the record medium and the program repeats the process with the next D record in the current row and thereafter the first D record in the next subsequent row.

In summary, the invention provides a file format which surpasses the prior art significantly in numerous respects. First null fields or varying length fields less than their maximum size are compressible. Still further storage of column attributes is ideal for numerous reasons. Attributes of each column are stored precisely once. All attributes (column name, column type, precision, length or maximum length, nullability, code

pages) can be stored and the format easily expandable for additional column attributes. Also, all relational types are supported and the format is easily expandable for new types.

Still further positive attributes of the invention include that due to the absence of forced conversions of database data, there are no data loss problems attributable to the format and processing time tends to be minimised when the data change is between like-systems. Moreover, code page identification fields allow for a user option to accept or reject data from a different codepage environment. Support for additional, application specific records provide further benefit such as the inclusion of additional application specific (A) records. Such records may be used to store indexes of tables for example and further may permit multi-diskette storage of a single PC/IXF file.

Although a specific embodiment of the invention has been described above, the invention is not limited thereto and it will be appreciated that modification and/or additions thereto are possible within the scope of the appended claims.

Claims

1. A relational data table for a computer system comprising a plurality of items of data arranged logically at the intersections of a plurality of table rows and a plurality of table columns, the table being formed by a plurality of records, such records comprising;

(1) a group of one or more data ("D") records for each table row, each D record being of variable length and including one or more items of data for that row arranged serially in succession and terminating with a data character which may be null or of a variable length, the data items for the rows being ordered in a consistent manner within their respective groups of D records,

(2) a plurality of column descriptor ("C") records, each C record being associated with one table column and identifying the location of the record of each group which contains the data item for the column, and the beginning position of the associated data item column in the associated D record; whereby null or varying length data items which terminate a D record are compressible.

2. A relational data table as claimed in claim 1, wherein each C record includes an indication of the maximum number of characters forming an item of data in the associated column, if that number is not implied by the type of data in the data item.

3. A relational data table as claimed in claim 1 or claim 2, wherein the records are arranged in series with all of the C records preceding the D records.

4. A relational data table as claimed in any preceding claim, wherein a header ("H") record is included which H record specifies some or all of the following information:

- (i) the software program that generated the table;
- (ii) the length of the H record;
- (iii) the codepage environment of the table;
- (iv) the date and time of writing of the table.

5. A relational data table as claimed in any preceding claim, which further includes a table ("T") record which specifies one or more of the following items of information:

- (i) the type of computer with which machine format the data within the table is compatible;
- (ii) the length of the T record;
- (iii) the number of C records in the table;
- (iv) the name for the data contained in the table.

6. A relational data table as claimed in any preceding claim, wherein each D record further specifies its own length.

7. A relational data table as claimed in any preceding claim, wherein each D record includes a code which identifies such D record, and wherein the identification codes for the D records of any given row in the table are different from each other and in ascending order.

8. A relational data table as claimed in any preceding claim, wherein the identification of the associated D record in the C record includes the identification code.

9. A relational data table as claimed in any preceding claim, wherein each C record further specifies its own length.

10. A relational data table as claimed in any preceding claim, wherein the column identification of each C record includes a name for the associated column.

11. A relational data table as claimed in any preceding claim, wherein the C record further specifies whether nulls are permitted to be entered in the associated column.

12. A relational data table as claimed in any preceding claim, wherein the C record specifies the codepage environment of the data to be stored in the associated column.

5 13. A relational data table as claimed in any preceding claim, wherein the C record specifies the type of data which may be entered in the associated column.

14. A relational data table as claimed in any preceding claim, wherein the C record further specifies the length of the data item which may be entered in the associated column, if such length is not implied by the data type, the data length being the maximum length of the data item if the data item which may be entered
10 in the associated column is of variable length.

15. A relational data table as claimed in any preceding claim, wherein the C record includes a number specifying the position within a data record of the data item associated with the corresponding column.

16. A digital computer system maintaining on a record medium thereof a relational data table as claimed in any preceding claim.

15 17. A method of storing for accessing and/or transmitting, a relational data table of a digital computer system comprising a plurality of items of data logically arranged at the intersections of a plurality of table rows and a plurality of table columns, the method including the step of storing, on a record medium, the table as a plurality of records arranged successively in series, such records comprising;

(1) a group of one or more data ("D") records for each table row, each D record being of variable length
20 and including one or more items of data for that row arranged serially in succession and terminating with a data character which may be null or of a variable length, the data items for the rows being ordered in a consistent manner within their respective groups of D records,

(2) a plurality of column descriptor ("C") records, each C record being associated with one table column and identifying the location of the record of each group which contains the data item for the column, and
25 the beginning position of the associated data item column in the associated D record;
whereby null or varying length data items which terminate a D record are compressible.

30

35

40

45

50

55

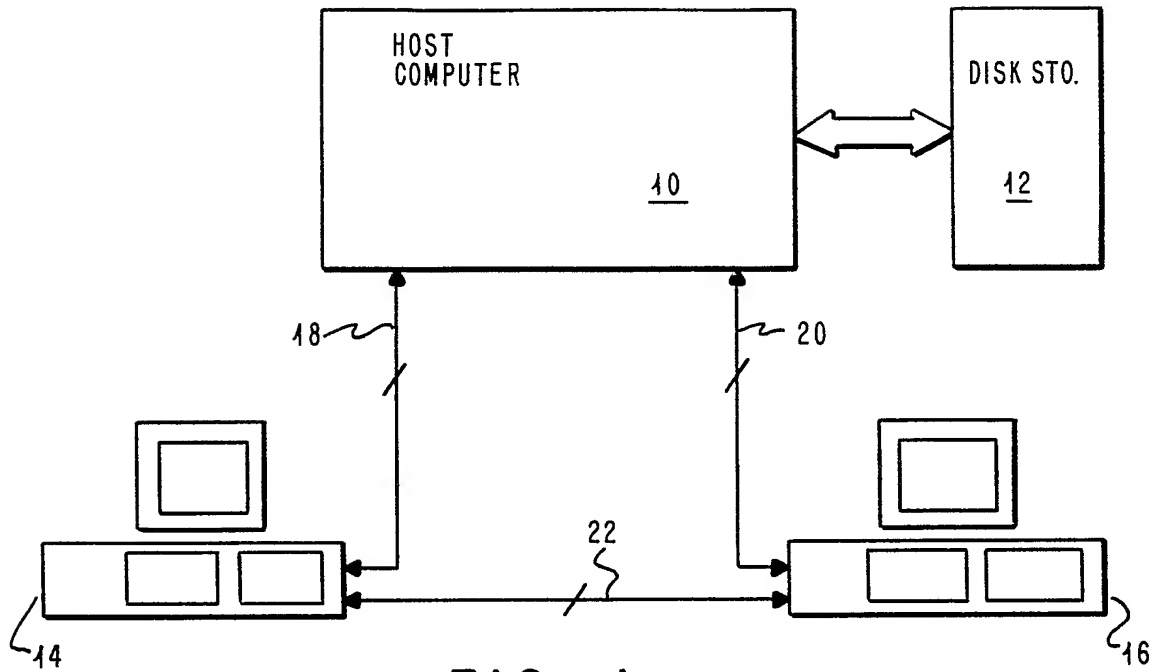


FIG. 1

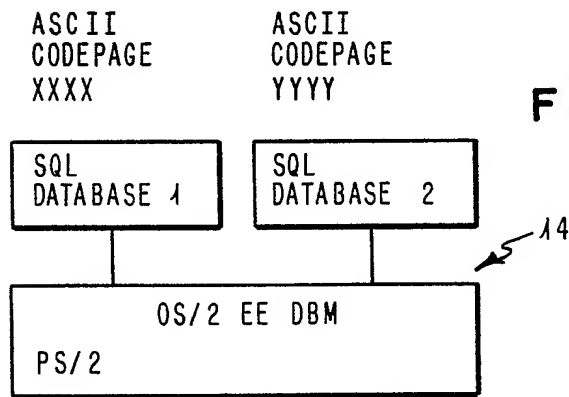


FIG. 2

FIG. 3

TABLES - CITIES

	NAME	STATE	REGION	POPULATION
→	AUSTIN	TX	SW	450000
→	CHICAGO	IL	MW	3200000
→	DALLAS	TX	SW	1700000
→	SAN ANTONIO	TX	SW	1100000
→	SAN JOSE	CA	W	700000

ROWS

ONE PER CITY

COLUMNS - ONE PER INFORMATION FIELD

COLUMN NAME	NAME	STATE	REGION	POPULATION
DATA TYPE # CHARACTERS NULLABLE ?	VARCHAR 30 REQUIRED	CHAR 2 REQUIRED	VARCHAR 30 NULLABLE	INTEGER 8 REQUIRED
D RECORD NUMBER	1	2	2	3
D RECORD POSITION	BEGINNING (OF DATA PORTION OF D RECORD)	BEGINNING	BEGINNING PLUS LENGTH OF STATE COLUMN	BEGINNING

FIG. 4

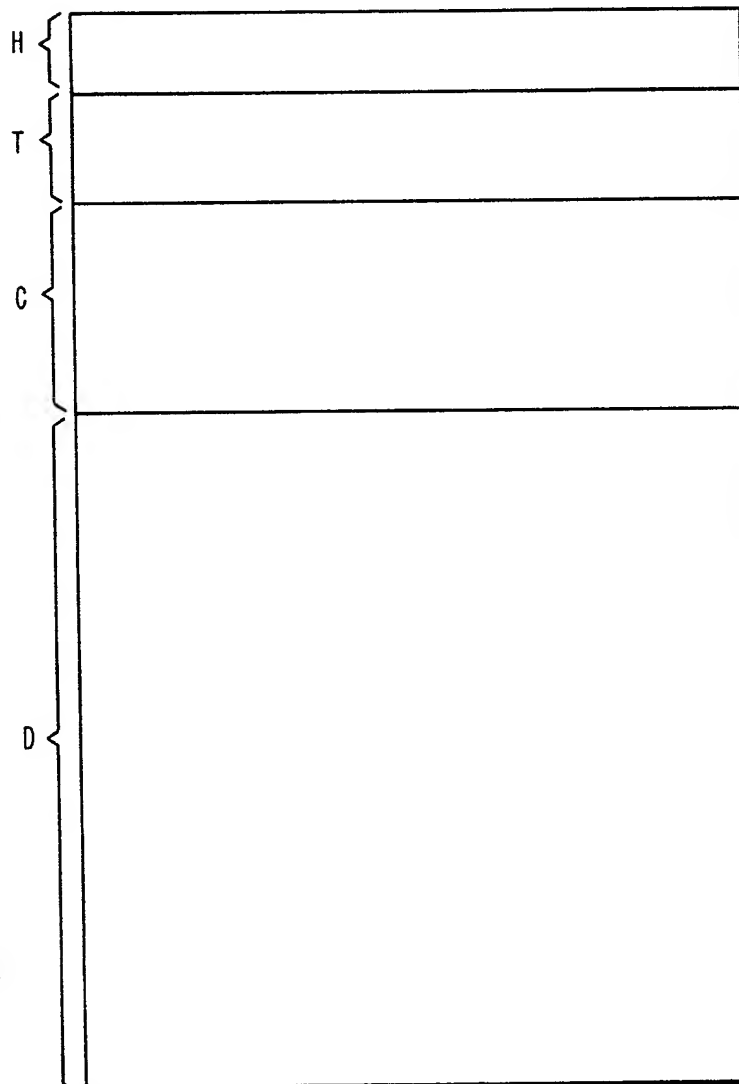


FIG. 5

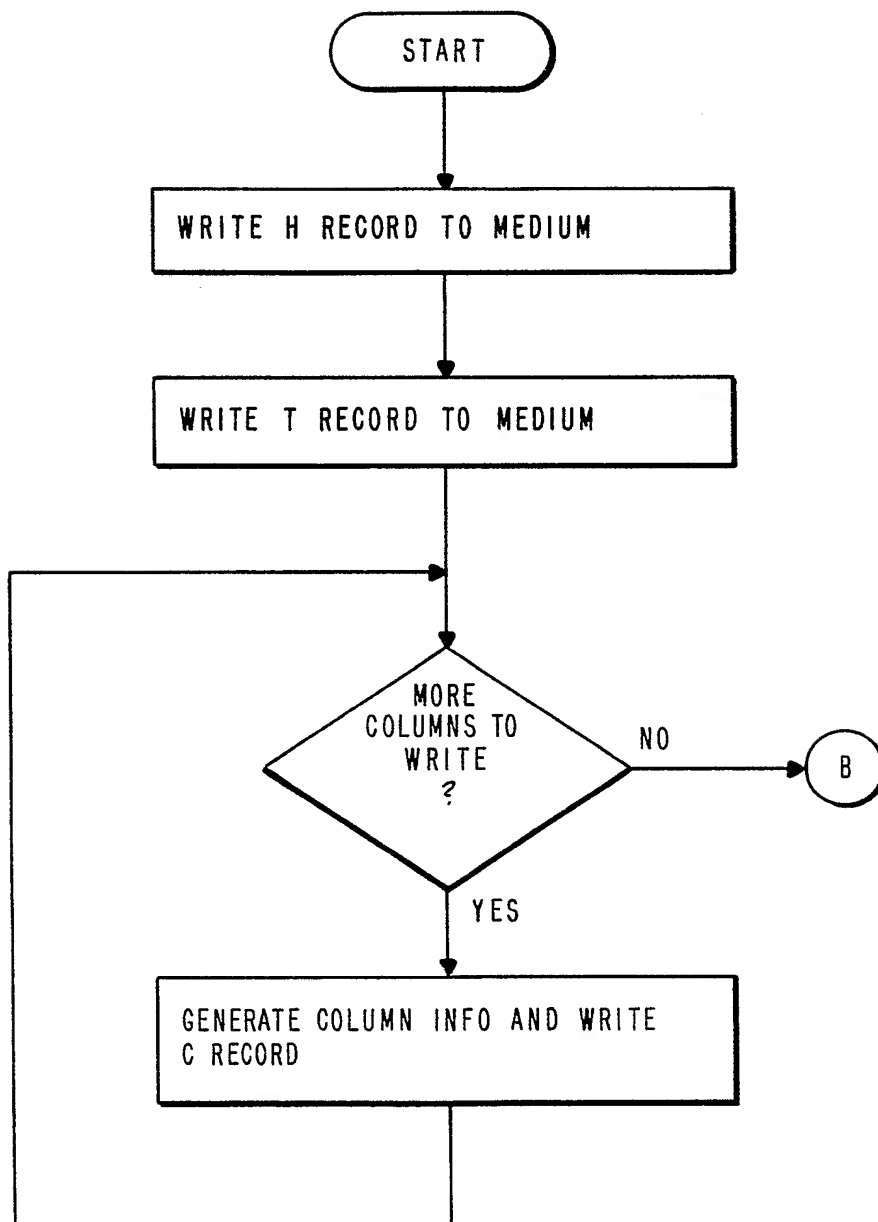


FIG. 6A

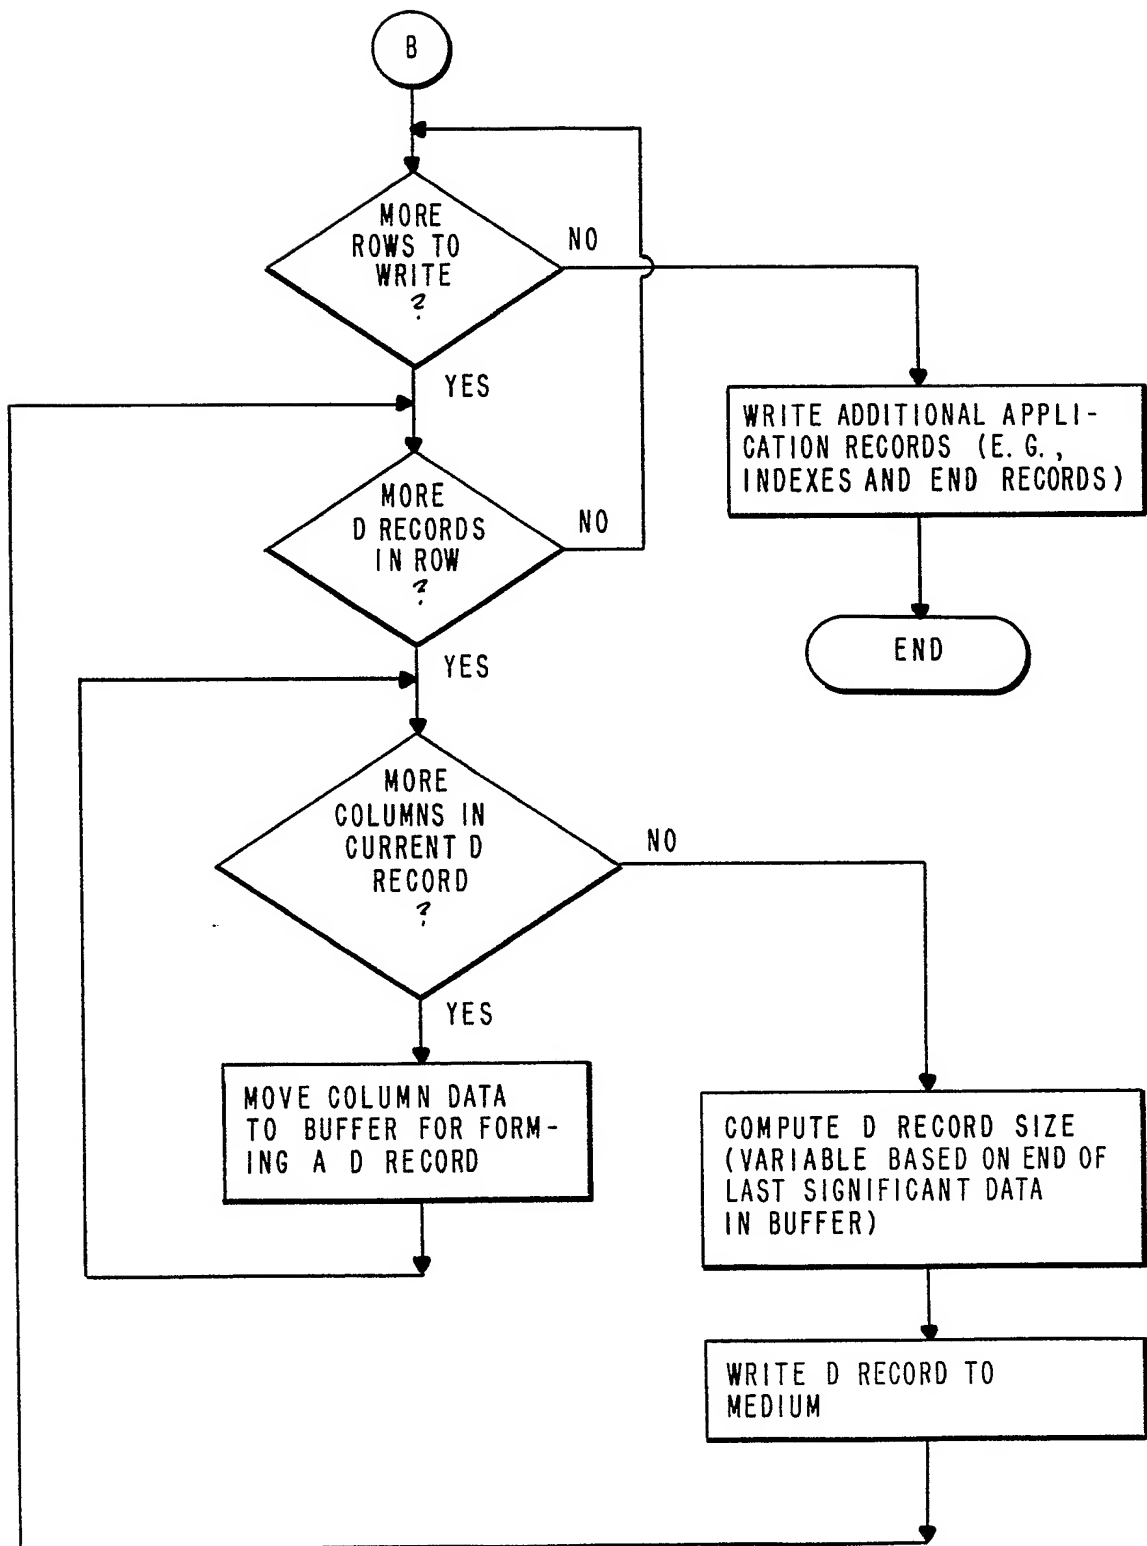


FIG. 6B